Random Lifts of Graphs

Alon Amit *
alona@math.huji.ac.il

Nathan Linial †
nati@cs.huji.ac.il

Jiří Matoušek ‡
matousek@kam-enterprise.ms.mff.cuni.cz

Eyal Rozenman §
eyalroz@cs.huji.ac.il

July 7, 2000

Abstract

We describe here a simple probabilistic model for graphs that are lifts of a fixed base graph G, i.e., those graphs from which there is a covering map onto G. Our aim is to investigate the properties of typical graphs in this class. In particular, we show that almost every lift of G is $\delta(G)$-connected where $\delta(G)$ is the minimal degree of G. We calculate the typical edge expansion of lifts of the bouquet B_d and

*Institute of Mathematics, the Hebrew University, Jerusalem, Israel.

†Institute of Computer Science, the Hebrew University, Jerusalem, Israel. Supported in part by a grant from the Israel academy of sciences.

‡Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic. Supported by Charles University grants No. 158/99 and 159/99.

§Institute of Computer Science, the Hebrew University, Jerusalem, Israel.

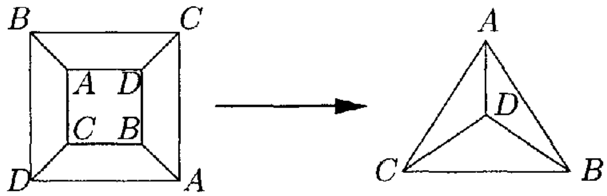

Figure 1: The cube is a lift of K_4

other base graphs. Combined with the novel method of ϵ-nets this allows us to construct d-regular graphs whose edge expansion (slightly) exceeds previously known bounds. For the independence number, upper and lower bounds are obtained as solutions to certain optimization problems on the base graph. For a graph G with chromatic number χ and fractional chromatic number χ_f, we show that the chromatic number of typical lifts is bounded from below by $const \cdot \sqrt{\chi / \log \chi}$ and also by $const \cdot \chi_f / \log^2 \chi_f$. We have examples of graphs where the chromatic number of the lift equals χ almost surely, and others where it is a.s. $O(\chi / \log \chi)$. Finally, we present conditions on the base graph G that determine whether lifts of G have perfect matchings surely, almost surely, or matchings that necessarily miss a logarithmic or linear number of vertices.

1 Introduction

We begin with an example of a covering map f from Q_3, the three dimensional cube to the complete graph K_4, shown in Figure 1. (Vertices in Q_3 are labeled according to their image under f.) This map is a local isomorphism, where the neighbors of every vertex x in Q_3 are mapped one-to-one onto the neighbors of $f(x)$. Such a mapping is called a *covering* (from Q_3 onto K_4), and we say that Q_3 is a *lift* of K_4. It is no coincidence that the inverse image of every vertex of K_4 has the same cardinality (here — two). The cardinality of the preimage of a vertex is called the *degree* or *fold number* of the lift.

Fix a *base graph* G — here K_4. Lifts of G, like Q_3 in our example, all share the local structure of G. But how about global properties? Q_3 is still 3-connected, but in chromatic number and the girth, for example, Q_3 and K_4 differ. More interestingly, when the fold number grows, we obtain a family of graphs sharing the local structure of the base graph, and it is interesting to investigate the global properties of such lifts.

Lifts and covering maps have received considerable attention from several different perspectives. Angluin [3] used them in the study of distributed networks. Another application to distributed computation appears in [13] where they are the basis of a new, conceptual proof of the Fischer-Lynch-Patterson Theorem [14]. Angluin has also raised the problem of characterizing pairs of graphs that share a common finite lift, a question which Leighton [10] resolved. Leighton's result turned to be crucial for deep algebraic investigations regarding groups acting on trees (Bass and Kulkarni [5]). For further connections with group theory see, e.g., Stillwell [17] and Stallings [16]. Finally, as purely graph theoretic objects, lifts were studied by Gross and Tucker [7], [8], Negami [15] and Archdeacon and Richter [4], among others.

The main new ingredient of this work is the introduction of a probability distribution on the n-lifts of a fixed base graph. The probabilistic method can thus be employed to study lifts, and specifically to investigate the properties of a *typical* lifts. In this extended abstract we deal with connectivity, edge expansion, independence number, chromatic number and matching properties.

As usual, one of the advantages of the probabilistic approach is that it allows us to prove the existence of graphs (in the present case, lifts) with properties that may be hard to achieve otherwise. Moreover, we can "forget" that the random graphs we generate are lifts some base graph, and view our model as some sort of random model for (a special class of) graphs. Inasmuch as the probabilistic method is used to prove the existence of graphs with desired properties, a new model for random graphs (complete with "knobs" we can fiddle with) may be a useful addition to our arsenal.

In a different direction, we may hope that the model can be used to "randomize" a given graph, in the following sense. Suppose we wish to calculate, or approximate, some invariant of a graph G. It is often the case that the invariant is hard to determine efficiently, but an efficient algorithm is known to work on a *random* graph. We can now run the algorithm on a random lift $\tilde{G}$ of G, hoping that it is "sufficiently random" so that the algorithm runs reliably, but still "sufficiently similar" to G itself so the result is close to the original invariant. The results on the chromatic number below may be a first step in that direction.

This abstract is based in three written papers and one paper in preperation ([1], [2], [11], [12]). No proofs are provided here - only sketches. For complete proofs the reader is referred to the original papers which may currently be viewed at `http://www.cs.huji.ac.il/~eyalroz`

2 The Model

Let G and $\tilde{G}$ be finite graphs and let $\pi : \tilde{G} \to G$ be a graph homomorphism, namely a map from $V(\tilde{G}) \to V(G)$ that preserves adjacency. π is a *covering* if the edges incident with v are mapped bijectively onto the edges incident with πv, for every v in $\tilde{G}$. Moreover, π is required to be onto the vertices of G, but this condition follows automatically when G is connected, as will always be the case here.

We call G the *base graph* and $\tilde{G}$ a *lift*[1] of G. The inverse images $\pi^{-1}(v)$ are the *fibers*, denoted $\tilde{G}_v$. It is easy to see that all the fibers have the same cardinality. This common cardinality is called the *degree* or *fold number* of the lift. An n-lift is a lift of degree n.

Given the base graph G and the degree n, we can construct a lift $\tilde{G}$ by putting n vertices $v_1, \ldots, v_n$ above every vertex v of G. For every edge $e = [u, v]$, we need to decide which u_i is connected to which v_j. This is determined by a single permutation on n elements: Given a permutation $\sigma = \sigma(e)$ we connect u_i to $v_{\sigma(i)}$, and this is done independently for every edge. Notice that we need to assign an orientation to e so that we know which end is u and which is v. This choice, however, is arbitrary and has no real effect on the possible outcomes.

It is now fairly obvious how to define a *random* n-lift: simply choose a permutation $\sigma(e) \in S_n$ uniformly and independently for every edge e in G, and form the lift $\tilde{G}$ as above. This is our model for a random n-lift of G (more precisely, this is a random *labeled* lift, as explained below). We summarize it in the following definition:

Definition 1 *Given a graph G, a* random labeled n-lift *of G is obtained by arbitrarily orienting the edges of G, choosing a permutation σ_e in S_n for each edge e uniformly and independently, and constructing the graph $\tilde{G}$ with n vertices $u_1, \ldots, u_n$ for each vertex u of G and edges $(u_i, v_{\sigma_e(i)})$ whenever $e = (u, v)$ is an oriented edge. A covering map $\pi : \tilde{G} \to G$ is defined by $\pi(u_i) = u$ and $\pi((u_i, v_j)) = (u, v)$.*

We note that the model actually gives a little more than what we required, since the resulting graph $\tilde{G}$ is equipped with a labeling $\{1, \ldots, n\}$

[1]We use this term to avoid confusion with other uses of "covering" in graph theory.

of the vertices in each fiber. Such a lift $\tilde{G}$ is called a *labeled* lift, and it is these objects for which we have a random model. The situation is analogous to that of random graphs, where the standard model is defined for labeled graphs instead of (isomorphism classes of) abstract graphs. Under some mild assumptions on the base graph, asymptotic results in our model are valid for the unlabeled case as well:

Theorem 1 *Let G be a graph with $|E(G)| > |V(G)|$, and let P be a property of unlabeled lifts of G. Then almost every labeled lift has property P iff almost every unlabeled lift has it.*

3 Properties of Lifts

We now survey the results concerning graphical properties of lifts. In most cases we deal with properties of almost all lifts, namely we prove that a random lift has a certain property with probability tending to 1 as n, the degree of the lift, tends to ∞. In some cases, our results are tight, while in others interesting questions remain open.

3.1 Connectivity

Fix a base graph G. What is the typical connectivity of a random lift of G? If δ is the smallest vertex degree in G, then every lift of G contains vertices of degree δ and is, therefore, at most δ-connected. We prove that a random lift is indeed almost surely δ-connected when $\delta \geq 3$:

Theorem 1 *Let G be a connected simple graph with minimal degree $\delta \geq 3$. Then almost all n-lifts of G are δ-connected.*

The probabilistic part of the proof involves lifts of the multigraph $K_2(\delta)$, the graph with two vertices and δ edges. We analyze the connectivity of typical lifts of this graph which are, in a suitable sense, "almost" δ-connected. In the general case, we assume that the parts of the lift $\tilde{G}$ that lie above topological $K_2(\delta)$'s in G behave in the typical way, which enables us to show that most sets in $\tilde{G}$ have δ neighbors. Various technical considerations are required to handle all other sets. We remark that showing that almost every lift is $(\delta-1)$-connected is relatively easy. Considerably more work is required to prove the tight result.

3.2 Edge Expansion

In this section we discuss the following problem: Fix a base graph G. What can be said about the typical edge expansion of a random lift $\tilde{G}$ of G? We define edge expansion as usual:

Definition 2 *Let G be a graph with v vertices. For $S \subset V(G)$, Let ∂S be the set of edges with one vertex in S and one outside S. The* edge expansion *$\xi(S)$ is defined to be $|\partial S|/|S|$, and the edge expansion of G is*

$$\xi(G) = \min\{\xi(S) \mid S \subset V(G), |S| \leq v/2\}.$$

It is easy to see that a lift $\tilde{G}$ of G cannot have higher edge expansion than G. Our main results here are the following.

Theorem 2 *Let $G = (V, E)$ be a connected graph with $|E| > |V|$. Then there is a positive constant $\xi_0 = \xi_0(G)$ such that almost every lift of G has edge expansion at least ξ_0.*

The proof provides a constant ξ_0 explicitly, however it does not seem to be very tight. More precise information can be obtained for some specific base graphs, such as the bouquet B_l (the graph with 1 vertex and l loops):

Theorem 3 *Almost every lift $\tilde{G}$ of $G = B_l$ has edge expansion $\xi(\tilde{G}) \geq \xi_0$, where ξ_0 is the smaller root of $H(\xi/d) = (d-2)/d$, $d = 2l$.*

Similar estimates are available for complete graphs. Numerically, these results yield exactly the same bounds as Bollobas obtains for random regular graphs in [6]. Using the method of ϵ-nets, we are able to improve these bounds, through the following result.

Theorem 4 *For every $d \geq 3$ there is an $\epsilon_0 = \epsilon_0(d)$ such that if $0 \leq \epsilon \leq \epsilon_0$ then the following holds. Let μ_0 be the larger solution of*

$$\frac{2}{d}(1 - H(\epsilon)) = 1 - H(\mu_0)$$

then the edge expansion of a d-regular graph is almost surely at least $d(1 - \mu_0 - 2\epsilon)$.

3.3 Independence Number

Any fiber in a lift $\tilde{G}$ of a graph G without loops is an independent set by definition. Therefore large independent sets in $\tilde{G}$ can be obtained by taking a union of fibers over independent vertices in G. Interestingly, this yields the largest independent sets in some cases, while in others it is possible to get larger independent sets by using some of the vertices in adjacent fibers. For example, we have the following result.

Proposition 1 *Let $V(G) = \{v_1, \dots, v_r\}$ and suppose that $\{\alpha_1, \dots, \alpha_r\}$ satisfy for every $1 \leq k \leq r$*

$$0 \leq \alpha_k \leq \prod_{\substack{i<k \\ [v_i,v_k]\in E(G)}} (1-\alpha_i). \tag{1}$$

Let $S = \sum \alpha_i$. For every $\varepsilon > 0$, a random lift $\tilde{G}$ of G almost surely contains an independent set of size $n(S - \varepsilon)$.

To present an upper bound on the independence number $\alpha(\tilde{G})$ of typical lifts of G we need the following optimization problem. Given a function $f : V(G) \to [0,1]$, let

$$h(f) = \sum_{v\in V(G)} H(f(v)) - \sum_{[u,v]\in E(G)} I(f(u), f(v))$$

where H is the entropy function and $I(p,q) = H(p) + H(q) - H(p,q)$. Now let

$$\tilde{a}(G) = \max_f \left\{ \sum_v f(v) \,\middle|\, h(f) \geq 0 \right\}$$

We prove:

Theorem 5 *Almost every n-lift $\tilde{G} \to G$ satisfies*

$$n\alpha(G) \leq \alpha(\tilde{G}) \leq n\tilde{a}(G).$$

Together, these results enable us to essentially determine the independence number of lifts of the complete graph K_{r+1}: Almost surely, $\alpha(\widetilde{K_{r+1}}) = \Theta_r(n \log r)$.

3.4 Chromatic Number

As usual, the problem here is to determine the distibution of the chromatic number $\chi(\tilde{G})$ of random lifts of G. It is convenient to define:

Definition 3 *Given a graph G, let*

$$\tilde{\chi}_h(G) = \min\{k \mid \chi(\tilde{G}) \leq k \textit{ for a.e. lift } \tilde{G} \textit{ of } G\}$$

$$\tilde{\chi}_l(G) = \max\{k \mid \chi(\tilde{G}) \geq k \textit{ for a.e. lift } \tilde{G} \textit{ of } G\}$$

These are the essential upper and lower bounds on the chromatic number of lifts of G. For non-trivial graphs we have $2 \leq \tilde{\chi}_l(G) \leq \tilde{\chi}_h(G) \leq \chi(G)$. A natural conjecture is that the chromatic number of random lifts satisfies a zero/one law and is essentially single-valued:

Conjecture 1 *For every graph G, $\tilde{\chi}_l(G) = \tilde{\chi}_h(G)$.*

We know that this is true for various classes of graphs: bipartite graphs, cubic graphs and certain "blow-up" graphs defined below. We also suspect that $\tilde{\chi}_l(G) \geq C \cdot \chi(G)/\log\chi(G)$ for some absolute constant C. In this direction we prove:

Theorem 6 *For every graph G,*

$$\tilde{\chi}_l(G) \geq \sqrt{\frac{\chi(G)}{3\log\chi(G)}}$$

A better estimate is possible if we replace the chromatic number $\chi(G)$ by the fractional chromatic number $\chi_f(G)$, defined as the minimum total weight of a linear combination of independent sets, such that the weight at each vertex is at least 1.

Theorem 7 *For every graph G,*

$$\tilde{\chi}_l(G) \geq \Omega\left(\frac{\chi_f(G)}{\log^2\chi_f(G)}\right).$$

The proofs of these results utilize the upper bounds on independent sets mentioned above. A different approach, using a deep theorem of Kim [9], yields an *upper* bound on the chromatic number for graphs with large maximum degree:

Theorem 8 *Let G be a graph with maximal degree $\Delta = \Delta(G)$. Then*

$$\tilde{\chi}_h(G) \leq \frac{\Delta}{\ln \Delta}(1 + o_\Delta(1))$$

For complete graphs, we now have

Corollary 1 *There exist constants $A > B > 0$ such that*

$$A\frac{r}{\log r} \geq \tilde{\chi}_h(K_r) \geq \tilde{\chi}_l(K_r) \geq B\frac{r}{\log r}.$$

Thus the chromatic number of lifts of complete graphs drops from r to $r/\log r$. On the other hand, we have:

Proposition 2 *For any graph G with $\chi(G) \geq 2$, put $r = 3\chi(G)\log\chi(G)$, and let H be constructed from G by replacing each vertex by an independent set of size r and every edge by a complete bipartite graph $K_{r,r}$. Then almost every lift $\tilde{H}$ of H has chromatic number $\chi(\tilde{H}) = \chi(H) = \chi(G)$.*

3.5 Matchings

When the base graph G has a perfect matching (PM), so does every lift $\tilde{G}$ of G. However, it is possible that G does not have a PM while almost every lift does. Our aim is to characterize these cases, and determine how far from having a PM are the lifts of other base graphs.

When $V(G)$ is odd, n-lifts have an even or odd number of vertices depending on n, so we can never actually claim that almost every lift has a PM. To overcome this technical detail we define an *almost-perfect matching*, or APM, as a matching that misses at most one vertex. To define the class of graphs for which a.e. lift has an APM, we need the following definitions:

Definition 4 *A* fractional matching *in a graph $G = (V, E)$ is a mapping $f : E \to \mathbb{R}^+$ such that $\sum_{e=[v,x]} f(e) \leq 1$ for every vertex $v \in V$. If equality holds at every vertex, f is called a* perfect fractional matching *(PFM in short).*

Graphs having a PFM can be characterized similarly to Tutte's condition for PM.

Definition 5 *Let $I(H)$ be the number of isolated vertices in a graph H (A vertex with a loop is not considered isolated). The* excess *of a graph G is* $\min\{|S| - I(G \setminus S)\}$ *over all $S \subsetneq V(G)$ with $I(G \setminus S) > 0$.*

Lemma 1 *A graph $G = (V, E)$ has a PFM if and only if it has a nonnegative excess.*

Our basic result on matchings in lifts is

Theorem 9 *Let G be a graph that satisfies the following conditions:*

1. *G is connected.*
2. *$|E(G)| > |V(G)|$.*
3. *G has excess at least 1 (In particular, G has a PFM).*

Then almost every lift of G has an APM.

A PM in a lift naturally determines a PFM in G, by considering the proportion of edges in the matching that lie above a given edge of G. This explains why a PFM is a necessary condition for a PM in lifts to be even possible. The proof of the theorem involves various reduction operations, and direct handling of the minimal cases which are "loop stars", graphs obtained from $K(1, t)$ by adding a loop at each leaf.

The general classification result is:

Theorem 10 *Let G be finite connected graph. Exactly one of the following situations occurs:*

1. *Every lift $\tilde{G}$ of G has a PM. This occurs when G has a PM.*
2. *Almost every lift $\tilde{G}$ of G has an APM (theorem 9).*
3. *In almost every n-lift, the largest matching misses $\Theta(\log n)$ vertices. This happens e.g. when G is an odd cycle.*
4. *Every matching in every n-lift misses $\Omega(n)$ vertices. This happens if $\sum f(e) \leq (1/2 - \epsilon)|V|$ for every fractional mathcing in G.*

The implicit constants in the Θ and Ω terms depend only on G.

References

[1] Alon Amit and Nathan Linial. Random lifts of graph I: General theory and connectivity. To appear.

[2] Alon Amit and Nathan Linial. Random lifts of graphs II: Edge-expansion. To appear.

[3] Dana Angluin. Local and global properties in networks of processors. In *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*, Los Angeles, California, 1980.

[4] Dan Archdeacon and R. Bruce Richter. On the parity of planar covers. *Journal of Graph Theory*, 14(2), 1990.

[5] Hyman Bass and Ravi Kulkarni. Uniform tree lattices. *Journal of the Americn Mathematical Society*, 3(4):843–902, October 1990.

[6] Béla Bollobás. The isoperimetric number of random regular graphs. *European Journal of Combinatorics*, 9:241–244, 1988.

[7] J. L Gross and T. W. Tucker. Generating all graph coverings by permutation voltage assignments. *Discrete Mathematics*, 18:273–283, 1977.

[8] J. L. Gross and T. W. Tucker. *Topological Graph Theory.* Wiley-Interscience, 1987.

[9] Jeong Han Kim. On Brooks' theorem for sparse graphs. *Combin. Probab. Comput.*, 4(2):97–132, 1995.

[10] Frank Thomson Leighton. Finite common coverings of graphs. *Journal of Combinatorial Theory Ser. B*, 33:231–238, 1982.

[11] Alon Amit Nathan Linial and Jiří Matoušek. Random lifts of graphs III: Independence and chromatic number. To appear.

[12] Nathan Linial and Eyal Rozenman. Random lifts of graphs IV: Matching. To appear.

[13] N. A. Lynch M. J. Fischer and M. Merritt. Easy impossibility proofs for distributed consensus problems. *Distributed Computing*, 1:26–39, 1986.

[14] N. A. Lynch M. J. Fischer and M. Patterson. Impossibility of distributed concensus with one faulty processor. *J. ACM*, 32:374–382, 1985.

[15] Seiya Negami. The spherical genus and virtually planar graphs. *Discrete Mathematics*, 70:159–168, 1988.

[16] John Stallings. Topology of finite graphs. *Inventiones Mathematicae*, 71:551–565, 1983.

[17] John Stillwell. *Classical Topology and Combinatorial Group Theory.* Springer-Verlag, 1993.